

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 162 790 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
12.12.2001 Bulletin 2001/50

(51) Int Cl.7: **H04L 12/56**

(21) Application number: **01304998.6**

(22) Date of filing: **08.06.2001**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR**
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **09.06.2000 US 210865 P
20.11.2000 US 716922**

(71) Applicant: **Hughes Electronics Corporation**
EI Segundo, California 90245 (US)

(72) Inventor: **Barnhart, Andrew**
Gaithersburg, Maryland 20882 (US)

(74) Representative: **Jackson, Richard Eric et al**
Carpmaels & Ransford,
43 Bloomsbury Square
London WC1A 2RA (GB)

(54) Queue and scheduled utilization metrics

(57) An approach for measuring scheduler efficiency is disclosed. A scheduler (603) has metrics measurement capability, in which scheduler counters (803, 805, 807, and 809) correspond to scheduling events. A counter logic (801) selectively increments the scheduler counters (803, 805, 807, and 809). The counter logic

(801) is coupled to queues (607) that store cells. The performance metrics are computed based upon values of the counters (803, 805, 807, and 809). The scheduler (603), according to one embodiment of the present invention, is deployed in a satellite communication system.

DescriptionCROSS-REFERENCES TO RELATED APPLICATION

5 [0001] This application is related to, and claims the benefit of the earlier filing date of, U.S. Provisional Patent Application (Application Number 60/210,865), filed June 9, 2000, entitled "Scheduler Utilization Metrics," the entirety of which is incorporated herein by reference. Additionally, this application is related to U.S. Patent Application (Application Number 60/210,864), filed June 9, 2000, and entitled "Available Bandwidth Control Mechanism," which are incorporated herein by reference.

BACKGROUND OF THE INVENTIONField of the Invention:

15 [0002] The present invention relates generally to communication systems, and is more particularly related to providing congestion avoidance in a switching communication system.

Discussion of the Background

20 [0003] As businesses and society, in general, become increasingly reliant on communication networks to conduct a variety of activities, ranging from business transactions to personal entertainment, these communication networks continue to experience greater and greater traffic congestion. For example, the maturity of electronic commerce and acceptance of the Internet as a daily tool pose an enormous challenge to communication engineers to develop techniques to reduce network latency and user response times. With the advances in processing power of desktop computers, the average user has grown accustomed to sophisticated multimedia applications, which place tremendous strain on network resources (e.g., switch capacity). Also, because the decrease in application response times is a direct result of the increased processor performance, the user has grown less tolerant of network delays, demanding comparable improvements in the network infrastructure.

30 [0004] Traffic control can be accomplished using two general approaches: flow control, and congestion control. Flow control seeks to regulate the amount of traffic that is transmitted from a source station to a destination station, by permitting the destination station to control the rate at which the source transfers data as to not overload the respective destination node. Flow control, however, does not directly address the problems associated with managing the traffic load on the network; for instance, numerous source stations can be communicating at rates that are acceptable to the destination stations. Nonetheless, the network (i.e., networking components) may not be able to sustain the overall level of traffic that is exchanged by these source and destination stations. It should be noted that these flow-control protocols, as with Transfer Control Protocol (TCP), tend to cause network congestion - rather than avoid network congestion - by collectively driving the network until it exhibits packet loss along with maximum queuing.

35 [0005] Therefore, congestion avoidance schemes are needed to control the network traffic in a way as to effectively maintain the overall traffic that is introduced by the stations, which are generally connected via intervening nodes. Without a congestion avoidance scheme, a large queuing delay will occur, potentially resulting in dropped packets. Consequently, the quality of service of the system will likely be degraded. Also, because the switch is forced to drop packets, the useful capacity of the system is reduced; thus, the network service provider suffers a loss of revenue.

45 [0006] One traditional approach to congestion avoidance is to upgrade the hardware to increase capacity to enhance the throughput of the system. The main drawbacks with this forklift approach are cost and potential lack of interoperability. Further, in some systems, such as a communication satellite, hardware replacement is impractical. Furthermore, with wireless systems, additional frequency assignments may not be available. Another technique to avoid congestion involves the development of sophisticated networking protocols. One drawback with this approach is that the developed protocol may not be easily standardized; without industry acceptance, the development cost cannot be recouped. Another drawback is that the protocol is likely to be inefficient, requiring significant overhead bits to effect congestion control: this protocol inefficiency negatively impacts throughput of the network.

50 [0007] Based upon the forgoing, there is a clear need for improved approaches for the management of congestion to improve the traffic transmission efficiencies of networking components, and particularly, those components that are subject to potential traffic congestion.

[0008] Accordingly, it is highly desirable for the network to control the traffic that enters the network to avoid congestion and to maximize the effective network throughput without introducing excessive protocol overhead cost.

55 [0009] Congestion avoidance optimizing switching system performance is highly desirable. In this regard, both the metrics that are used in monitoring the traffic flow as well as the algorithms that are used to process these metrics are of significant interest.

SUMMARY OF THE INVENTION

[0010] According to one aspect of the invention, a method is provided for measuring scheduler efficiency. The method includes selectively incrementing a plurality of scheduler counters corresponding to a plurality of scheduling events. The method also includes computing performance metrics based upon values of the plurality of scheduler counters. Under this approach, congestion avoidance is provided, while increasing effective throughput.

[0011] According to another aspect of the invention, a scheduling device having metrics measurement capability comprises a plurality of scheduler counters corresponding to a plurality of scheduling events. A counter logic is configured to selectively increment the plurality of scheduler counters; the counter logic is coupled to a plurality of queues that store packets. The performance metrics are computed based upon values of the plurality of scheduler counters. The above arrangement advantageously enhances system efficiency.

[0012] In yet another aspect of the invention, a switching system comprises a plurality of input ports that are configured to receive packets. A plurality of output ports communicate with a plurality of destination sites. A plurality of queues are coupled to the input ports and the output ports for storing the received packets. A scheduler is coupled to the plurality of queues and is configured to assign the stored packets for transmission out the output ports. The scheduler comprises a plurality of scheduler counters corresponding to a plurality of scheduling events, and a counter logic that is configured to selectively increment the plurality of scheduler counters. A processor is configured to compute performance metrics based upon values of the plurality of scheduler counters. Under the above arrangement, a dynamically adaptive congestion mechanism is provided.

[0013] In yet another aspect of the invention, a computer-readable medium carrying one or more sequences of one or more instructions for measuring scheduler efficiency is disclosed. The one or more sequences of one or more instructions include instructions which, when executed by one or more processors, cause the one or more processors to perform the step of selectively incrementing a plurality of scheduler counters corresponding to a plurality of scheduling events. Another step includes computing performance metrics based upon values of the plurality of scheduler counters. This approach advantageously improves efficiency of a switching communication system that has transmission constraints.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] A more complete appreciation of the invention and many of the attendant advantages thereof will be readily obtained as the same becomes better understood by reference to the following detailed description when considered in connection with the accompanying drawings, wherein:

Figure 1 is a block diagram of a switch with congestion avoidance capability, in accordance with an embodiment of the present invention;

Figure 2 is a diagram of a communication network that utilizes the switch of Figure 1;

Figures 3A and 3B are diagrams of a satellite communication system with congestion avoidance capability, according to an embodiment of the present invention;

Figure 4 is a diagram of an interference region of a target downlink cell, in accordance with an embodiment of the present invention;

Figure 5 is a diagram showing the interaction among the satellite terminals (STs), the satellite, and the Network Operations Center (NOC) in a satellite communication system, in accordance with an embodiment of the present invention;

Figure 6 is a diagram of an available bandwidth control mechanism (ABCM) utilized in the system of Figure 5;

Figure 7 is a flow chart of the operation of the available bandwidth control mechanism of Figure 6;

Figure 8 is a diagram of the scheduler of Figure 5 utilizing four counters;

Figure 9 is a flowchart of the scheduler monitoring operation, according to an embodiment of the present invention;

Figure 10 is a flowchart of the metric collection process using the counters of Figure 8, according to an embodiment of the present invention; and

Figure 11 is a diagram of a computer system that can perform the functions of the ABCM to avoid congestion, in accordance with an embodiment of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0015] In the following description, for the purposes of explanation, specific details are set forth in order to provide a thorough understanding of the invention. However, it will be apparent that the invention may be practiced without these specific details. In some instances, well-known structures and devices are depicted in block diagram form in order to avoid unnecessarily obscuring the invention.

[0016] The present invention accomplishes congestion avoidance by using a scheduler that collects bandwidth metrics and transmits these metrics to a traffic control processing logic. A scheduler has metrics measurement capability, in which scheduler counters correspond to scheduling events. A counter logic selectively increments the scheduler counters. The counter logic is coupled to queues that store packets. The performance metrics are computed based upon values of the counters. The above approach provides for an enhanced bandwidth allocation scheme. The input to these queues may also be monitored by counters to further enhance the transient performance of the control system.

[0017] Figure 1 shows a functional block diagram of a switch with congestion avoidance capability, according to an embodiment of the present invention. Switch 100 includes multiple input ports 101 that receive incoming traffic from one or more source nodes (not shown) and forwards the traffic to a bandwidth control processor 103. A congestion avoidance logic 105 within switch 100 operates in conjunction with the bandwidth control processor 103 to control the amount of traffic entering input ports 101. A packet buffer 107 stores packets from input ports 101 that have been accepted for transmission by the bandwidth control processor 103. The stored packets in packet buffer 107 are transmitted to a scheduler 109, which communicates with a constraint check logic 111 to determine whether the stored packets conform with established transmission constraints. Scheduler 109 examines the destination addresses of the packets that are stored in packet buffer 107 using a round-robin scheme and forwards such packets to an appropriate output port among the various output ports 113. The congestion avoidance functionality of switch 100 is described below with respect to Figure 2.

[0018] Figure 2 shows a communication network that utilizes a switch with congestion avoidance functionality, in accordance with an embodiment of the present invention. A communication network 200 includes multiple source stations 201 that generate traffic to node 203, which can be any networking equipment that transfers data. In an exemplary embodiment, node 203 is an internetworking device, such as a router; alternatively, node 203 may be any type of gateway in a land-based or satellite-based communication system. Node 203 is connected to an input port (Figure 1) of switch 100. Although not shown, additional nodes, such as node 203, may be connected to additional ports 101 of switch 100. The output ports (Figure 1) of switch 100 connects to multiple nodes 205, which can be the same networking component as that of node 203. As shown, each of the nodes 205 can potentially communicate with numerous destination stations 207 within region 209 (e.g., sub-network). For example, if nodes 205 are routers, the routers would have multiple output ports designated for region 209.

[0019] As indicated previously, the conventional communication network exhibits performance characteristics that are dictated largely by the hardware limitations of switch 100. In other words, the throughput of the network 200 depend on such parameters as buffer size and processing capability of switch 100. In some practical systems, however, the communication network 200 possess network bottlenecks at points in the network other than the switch 100. For example, assuming that for security reasons, destination stations 207 within region 209 cannot simultaneously receive packets, consequently switch 100 may need to buffer some of the packets until the first set of packets are delivered to the particular destination stations 207. Thereafter, the buffered packets within switch 100 can be delivered to the destination stations 207 within region 209. From this example, it is clear that the buffering of the packets within switch 100 can result in system performance that does not depend on the hardware capabilities of switch 100, but instead on the network constraints associated with region 209.

[0020] The above scenario is characteristic of a satellite communication system. For the purposes of explanation, the operation of congestion avoidance according to an embodiment of the present invention, is discussed with respect to a satellite communication system with transmission constraints to downlink cells. However, it should be noted that the approach has applicability to many other switching systems, as recognized by one of ordinary skill in the art. For example, the switching systems may include an ATM (Asynchronous Transfer Mode)/SONET (Synchronous Optical Network) network, a Gigabit Ethernet network, and a voice network. The end stations of these systems are referred to as destination sites. Accordingly, the destination sites in a satellite system would be downlink cells.

[0021] Figures 3A and 3B show a satellite communication system with an on-board switch, in accordance with an embodiment of the present invention. The satellite payload 301 has a switch 302 that is connected to multiple transmitters 303; that is, N transmitters. Switch 302 includes a congestion avoidance logic 304, a bandwidth control processor (BCP) 305, a packet buffer 307, a downlink scheduler 309, and a constraint check logic 311. One of ordinary skill in the art would recognize that the congestion avoidance logic 304, bandwidth control processor (BCP) 305, downlink scheduler 309, and constraint check logic 311 may be implemented via software, hardware (e.g., general processor, an Application Specific Integrated Circuit (ASIC), etc.), firmware or a combination thereof. As shown, the BCP 305 is a part of switch 302. Alternatively, the BCP 305 may be implemented as a separate processor that is separate from switch 302 (as shown in Figure 6).

[0022] In switching communication system 300, satellite terminals (ST) (not shown) originate traffic from a coverage area 315. The generated traffic from the STs is transferred through switch 302 and terminate at destination STs (not shown) within coverage area 317. It should be noted that the destination STs can be within the same coverage area 315 as the originating STs. To effectively transmit traffic to the desired destination ST through switch 302, source STs transmit bandwidth requests to the bandwidth control processor 305 prior to transmitting any data traffic.

[0023] A connection that is established between a source ST and a destination ST is controlled by the BCP 305 and a network operation center (NOC). The NOC (shown in Figure 5), which is based on the ground, provides management functions for the system 300. An ST needs to obtain authorization from the NOC before making a request to the BCP 305. However, once an ST has received authorization from the NOC, the ST is likely to receive a rate allocation from the BCP 305 because the NOC keeps track of the total uplink (and downlink) bandwidth available for connections and will block a connection request if there is insufficient satellite capacity available to satisfy the request.

[0024] A major advantage of a congestion avoidance mechanism of system 300 is that if a connection request is blocked because of impending congestion in the switch 302, the connection may be given an alternate path for connection, should one exist, rather than admitting the connection to the switch 302 and potentially deteriorating the grade of service for both the connection and previously admitted connections. Consequently, the congestion avoidance mechanism of system 300 allows commitment to a quality of service for those connections that are admitted. A key source of congestion in the system 300 lies with the downlink channels 321. Accordingly, the bandwidth control processor (BCP) 305 implements the bandwidth control function which includes controlling the allocation of uplink channels and timeslots and mitigating downlink congestion. Bandwidth control processor 305 examines the requested bandwidth and replies with grants based on downlink resource availability, as determined by congestion avoidance logic 304 and uplink resource availability. As will be explained in greater detail later, congestion avoidance in the system 300 is achieved by the collaboration of bandwidth control processor 305, congestion avoidance logic 304, a scheduler 309, and a traffic control processing logic (613 of Figure 6) to regulate the amount of traffic received by the switch 302 through TDMA (Time Division Multiple Access)/FDMA (Frequency Division Multiple Access) uplink channels 319 via request/grant bandwidth control processes. The present invention concentrates on the interaction between the scheduler 309 and the traffic control processing logic, which are more fully described in Figures 5-7.

[0025] The service areas 315 and 317 are covered by a set of polygons that are fixed on the surface of the earth. Downlink polygons, called microcells, are hexagonal in shape as viewed from the spacecraft, with seven microcells clustered together to form an uplink polygon, called a cell. As used herein, the term microcell is used synonymously with the term downlink cell. The satellite generates a set of uplink circular beams that each encloses a cell. It also generates a set of downlink beams that each encloses a microcell.

[0026] Up to 12 downlink spot beams can be transmitted simultaneously on each of two polarizations subject to minimum microcell separation distance limitations. Beams on the same polarization must be sufficiently separated spatially to avoid unacceptable co-channel interference. Another co-polarized beam is not allowed to transmit to another microcell within an ellipse or else excessive interference may occur. The "keep-out" areas apply separately and independently for the two polarizations; the link budgets account for any cross-polarization interference that may occur.

[0027] As shown in Figures 3A and 3B, on the downlink of communication system 300, at each TDMA transmission slot, the downlink scheduler 309 selects up to n bursts of packets from M virtual queues of the packet buffer 307 to transmit through n transmitters, based on the scheduling algorithm and transmission constraint checks. The scheduling algorithm, in an exemplary embodiment, is a round-robin scheme. Sometimes the downlink scheduler 309 may not be able to find n bursts to transmit due to transmission constraints, in which case aggregate downlink transmission capacity will be limited by the transmission constraints. The downlink congestion in communication system 300 occurs when the amount of traffic admitted to the switch 302 exceeds the capacity of the downlink. In other words, if the BCP 305 made uplink allocations simply based on the availability of uplink slots, the BCP 305 would sometimes admit more traffic to a particular downlink cell (i.e., destination site) or cluster of mutually-interfering microcells than the downlink can carry. Consequently, the data packets for these areas would completely fill the packet buffer 307 in the payload's switch 302, resulting in dropped packets. Therefore, the availability of both uplink slots and downlink bandwidth factor into bandwidth allocations that is performed by the BCP 305.

[0028] The main transmission constraint in communication system 300 is the interference constraint; that is, two simultaneous downlink transmissions cannot be performed if they are directed at downlink cells which are within a system limit interference distance. In Figure 3B, because downlink cells A, B and C in coverage area 317 are outside the system limit interference distances (as shown by the overlapping circles) from one another, the satellite can simultaneously transmit packets to these downlink cells A, B and C. However, simultaneous transmission cannot be directed to downlink cells D and A, downlink cell D and B, downlink cell E and B and downlink cell E and C since they are within the system limit interference distance. That is, these downlink cells are in the same circle.

[0029] To illustrate the downlink capacity limitation of the system 300 stemming from transmission constraints (in particular, interference constraints), a scenario in which there are packets in the virtual queues of packet buffer 307 that are destined only to downlink cells A, B, C, D and E is considered. It should be noted that without any constraint, there can be five transmissions at one TDMA slot; however, with the interference constraint, only a maximum of three transmissions at one TDMA slot is possible (to downlink cells A, B, and C). If the satellite is to transmit to downlink cell D or downlink cell E, there can be at most two transmissions only, either to downlink cells D and C or to downlink cells E and A.

[0030] Figure 4 shows a diagram of an interference region of a target downlink cell defined according to an embod-

iment of the present invention. An interference region 400 includes a target downlink cell, which is surrounded by numerous downlink cells 401. Downlink cells 401 are clustered around target downlink cell Y within a radius that is determined by an angle x from the point of view of a satellite. The angle x can be set to any degree, depending on the coverage area and network application.

[0031] As will be discussed in more detail below, the BCP 305 and the congestion avoidance logic 304 limits the aggregate traffic going to a sets of downlink cells, referred to as an "interference cluster", instead of limiting the traffic going to each individual downlink cells. An interference cluster is a maximal set of downlink cells that are within the system limit interference distance from one another. Since clusters are not mutually exclusive of one another, a downlink cell can belong to more than one cluster.

[0032] According to one embodiment of the present invention, two types of requests for bandwidth allocation are defined: rate requests, and volume requests. In general, rate requests are utilized for connection-oriented traffic, while volume requests are used to transmit bursty traffic. In particular, rate requests specify the number of slots in each uplink frame that an ST needs to meet the uplink demands for a relatively constant traffic (e.g., connection-oriented). A rate request results in the allocation of a constant number of slots each frame, spread out as evenly in time as possible, which the ST can use to send packets at a constant rate. The requesting ST gets a constant allocation of that uplink capacity every frame until the request is cancelled by the ST via a de-allocation message to the satellite.

[0033] Volume requests specify the number of uplink slots that an ST requires to send a specific number of packets to another ST. The requesting ST receives a periodic allocation of zero, one or many slots within a specific frame until the entire number of slots requested has been allocated. Volume requests are used by the ST to send a burst (one or many) of data packets on the uplink. Several volume requests may be transmitted by the ST in a short period of time to send a file that has hundreds of data packets (e.g., segmented IP (Internet Protocol) packets) to another ST.

[0034] Figure 5 shows the interaction among the satellite terminals (STs), the satellite, and the Network Operations Center (NOC) in a satellite communication system, in accordance with an embodiment of the present invention. As shown, a satellite communication system 500 includes STs 501, 503, a satellite 505, and a NOC 507. For explanatory purposes, only the transmit section of a source ST and the receive section of a destination ST are shown. Accordingly, the ST 501 is a transmitting ST, while ST 503 is designated as a receiving ST. Each of the STs 501 and 503 has an interface 501b, 503b, respectively, for interfacing with an external network (not shown) and the satellite to control the traffic flow. Specifically, traffic from the external network (not shown) enters an input access 501b of the transmitting ST 501 and is transmitted to an uplink to the satellite 505. The satellite 505 forwards the traffic to the receiving ST 503 via the downlink. The traffic egresses an output access 503b of the receiving ST 503 to another network (not shown).

[0035] User traffic is received by the input access 501b of the transmitting ST 501, where it is stored and processed (e.g., segmentation of IP (Internet Protocol) frames). As previously discussed, for volume traffic the transmitting ST 501 makes bandwidth-on-demand (BoD) requests to a bandwidth control processor (BCP) 509. In turn, the BCP 509 selectively issues BoD grants that specify an uplink assignment, if it is determined that both uplink and downlink bandwidths are available. For rate traffic, the transmitting ST 501 makes traffic requests to the NOC 507, for which the NOC 507 authorizes according to the uplink and downlink bandwidths that are available. If so authorized, the transmitting ST 501 makes a rate BoD request to the BCP 509 to obtain a specific uplink assignment.

[0036] If transmitting ST 501 possesses a dedicated uplink assignment (i.e., rate request) such that the ST 501 does not employ BoD requests to the BCP 509 to send traffic, the ST 501 is considered a high-volume uplink (HVUL) ST. That is, an active HVUL ST 501 already has uplink bandwidth assigned by the NOC 507. For any significant change in traffic demand in HVUL traffic (to a specific destination), the ST 501 makes traffic rate adjustment request to the NOC 507, for which the NOC 507 provides authorization according to the downlink bandwidth that is available.

[0037] Traffic sent the uplink might be dropped in the satellite 505, if the switch queue 511 is experiencing congestion for a particular microcell. This operation is based upon both queue utilization and the priority level of the specific traffic involved. If the traffic is not dropped, the traffic will be queued in the switch queue 511 for transmission via a downlink to the microcell that contains the destination ST.

[0038] In the system 500 of Figure 5, every 21.7 μ sec, a downlink scheduler 513 assigns up to 24 bursts of traffic (e.g., each burst may be 12 packets) to the downlink spot beams. As previously discussed, a significant potential source of congestion is that the scheduler 513 needs to make these assignments to geographically disperse locations in order to avoid interference between neighboring microcells (i.e., transmission constrained, in part, because of interference considerations). If downlink congestion occurs, the queues 511 rapidly fill up and overflow, resulting in dropped packets.

[0039] The STs 501 and 503 support numerous application protocols, as represented by the input access 501b and output access 503b, in addition to utilizing set of networking protocols that are required to communicate with the external network (not shown). However, for switching efficiency, as with Asynchronous Transfer Mode technology, the network 500 employs a common packet size from the backbone interface (BBIF) 501a through the satellite 505 to the BBIF 503a. As a result, segmentation and reassembly (SAR) of the various protocol data units cannot be avoided. Effective throughput depends on the success of this SAR process, giving rise to the concept of "good" throughput ("goodput")

and "bad" throughput ("badput").

[0040] The effective throughput is the amount of traffic that exits the network and is useful to the end application. Goodput is a metric for evaluating this effective throughput. The minimum criteria for goodput is the number of successfully reassembled packets; goodput may also include end application effects as appropriate, such as TCP (Transmission Control Protocol) windows. For example, assuming that an IP packet is segmented into 15 segments at the transmitting ST 501, if all 15 segments are successfully recombined at the receiving ST 503, then all 15 segments contribute to goodput. However, if one segment is lost (e.g., during overflow of queues 511), then without a reliable link protocol (RLP), there is zero goodput for the IP packet.

[0041] Badput is that portion of the downlink traffic that is not useful to the end application, in particular that which does not get successfully recombined at the receiving ST 503. Continuing with the above example, if one segment is lost, then, without RLP, there are 14 segments of badput for the IP packet. With standard TCP/IP, the whole TCP window would constitute badput when a single segment is lost. The units for goodput and badput are the same as that used for throughput; further, these parameters may also be expressed as a percentage of the total throughput.

[0042] The switch (not shown) within the satellite 505 supports a packet drop priority scheme, in which the BCP 509 implements drop thresholds, which limit the number of packets entering the queues 511. In an exemplary embodiment, four drop priorities are implemented to determine the order in which packets will be dropped during congestion. Priority 0 is assigned to the highest priority traffic and would be the last type of packets to be dropped. Priority 3 is assigned to the lowest priority traffic and would be the first type of packets to be dropped. Before uplink transmission, the transmitting ST 501 marks packets as one of these four drop priorities.

[0043] The queues 511, which in an exemplary embodiment are logical, correspond to the microcells; the queues 511 store traffic on a burst basis (e.g., in groups of 12 packets). In an exemplary embodiment, for each one of the queues 511, a pair of thresholds (i.e., a drop maximum ("drop max") and a drop minimum ("drop min")), which correspond to each of the four priorities, restrict the number of packets that are stored in the queues 511. When the number of bursts in a queue 511 exceeds the drop max threshold for a given priority, the packets in the priority are dropped as they are demodulated from the uplink. Packets continue to be dropped until the traffic falls below the drop min threshold. If congestion persists, packets are dropped based upon the associated thresholds of the other priorities. Utilizing both a drop max and a drop min provides for a hysteresis in queue occupancy that "bunches" packet dropping. This capability is useful in reducing the generation of badput. Badput is also minimized by accounting for the utilization and available bandwidth of the microcells via the scheduler 513, in accordance with the present invention.

[0044] The scheduler 513 generates messages that contain utilization metrics to the NOC 507. According to one embodiment of the present invention, for each microcell, the scheduler 513 employs two counters (not shown) to gather bandwidth metrics associated with the utilization and available bandwidth of the microcell. These counters are denoted as follows: a "Yes" counter that tracks the number of times that an assignment was made for a specific microcell, and a "No" counter that tracks the number times a queue 511 was examined for assignment, but had no traffic to send. The counter values of both counters are sent to the NOC 507 after every measurement period; e.g., every super frame (786 msec).

[0045] The computation of the traffic statistics can be performed within the satellite 500, in which case the message that is sent to the NOC 507 would include these computed traffic statistics. Alternatively, a processor 508 within NOC 507 can compute the traffic statistics from the counter values that are received from the scheduler 513, as follows:

$$\text{Carried Load} = \text{Yes} \quad \text{Eq. (1)}$$

$$\text{Remaining Available Bandwidth} = \text{No} \quad \text{Eq. (2)}$$

$$\text{Total Available Bandwidth} = \text{Yes} + \text{No} \quad \text{Eq. (3)}$$

$$\text{Percent Utilization} = \text{Yes} / (\text{Yes} + \text{No}) \quad \text{Eq. (4)}$$

In Eq. (1), the Carried Load is equated with the Yes counter value. The Remaining Available Bandwidth, as seen in Eq. (2), equals the No counter value. The Total Available Bandwidth (Eq. (3)) is determined based upon the sum of the Yes and No counter values. The Percent Utilization, as seen in Eq. (4), is equal to the Yes counter value divided by the sum of the Yes and No counter values. It is observed that the utilization of Eq. 4 is responsive to both changing carried traffic, as given by the Yes counter value and to changing available bandwidth, as given by the Yes and No counter values. The NOC 507 processes the above bandwidth metrics and determines the bandwidth that is available

to each microcell. Then, according to service demands and policy rules, it apportions this bandwidth to the traffic classes: rate, HVUL, and volume.

[0046] For volume traffic, the NOC 507 determines a threshold for each microcell that the BCP 509 should not exceed in granting BoD requests for that destination during a bandwidth allocation period (96 msec.). These thresholds are set to minimize packet dropping, while not under utilizing available bandwidth. For each bandwidth allocation period, a BoD request for a microcell may be reduced or denied if making the grant with the others already granted would cause the traffic to exceed the threshold specified for the cell. It is noted that at the start of the bandwidth allocation period, pre-existing rate grants are immediately subtracted from the thresholds, with only the remainder being available for volume traffic.

[0047] An additional consideration arises from the fact that the BCP 509 has finite processing power. During periods of high traffic demand, the number of BoD requests to the BCP 509 may be significantly beyond the capacity of the BCP 509 to process them. Accordingly, a mechanism is provided for the transmitting ST 501 to throttle back the volume BoD requests to the BCP 509. The BCP 509 determines its utilization, which is periodically sent in control messages to the NOC 507, alerting the NOC 507 about the level of BCP congestion.

[0048] As a result, NOC 507 sends these control messages to the transmitting ST 501 to reduce ST BoD requests during periods of BCP congestion. The ST 501 may then be responsive to the reduction in its BoD requests - without potential loss of traffic - by simply aggregating more traffic for each request. This operation is generally easy to achieve during periods of high traffic demand. The transmitting ST 501 receives the BoD thresholds (an other bandwidth control parameters) from the NOC 507, and receives the bandwidth requests from the input access 501b. Then, for each cell that the ST 501 has traffic to send to, a BoD regulator 501c determines the actual BoD requests that are to be sent up to the BCP 509 for that cell. A BoD controller 501d within the transmitting ST 501 initiates the BoD requests to the BCP 509. Although shown separately, the BoD regulator 501c may be integrated with the BoD controller 501d.

[0049] Figure 6 shows a diagram of the available bandwidth control mechanism (ABCM) utilized in the system of Figure 5. A satellite communication system 600 includes numerous transmitting STs (not shown), which offer traffic 606 to a satellite 601. In turn, the satellite 601 forwards the traffic 608 to the appropriate receiving STs (not shown). It should be noted that all STs possess the capability to transmit and receive traffic (as mentioned previously); therefore, the labels, transmitting ST and receiving ST, pertain to the mode of operation of an ST at a particular time instance. The traffic reaches the receiving STs by downlink spot beams, which are electronically steerable by the satellite 601 to an earth surface area designated as a microcell.

[0050] Although possessing high bandwidth (for example, 450 Mbps), due to practical limitations of communications satellite design, the satellite 601 utilizes a relatively low number of spot beams when compared with the number of microcells. In an exemplary embodiment, the satellite communication system 600 has 1200 microcells, and the satellite supports 24 spot beams. Accordingly, the spot beams are time shared among the microcells. This is accomplished by a scheduler 603. The scheduling process is performed as required, in light of the fact that the receive traffic requirements of the various microcells may vary by a considerable amount, and that there may be competition among the microcells for available scheduler bandwidth. The scheduler 603 provides a self-monitoring function to output scheduler bandwidth metrics; for example, by utilizing the Yes and No counters, as discussed previously.

[0051] The satellite 601 provides temporary storage of traffic (i.e., packet bursts), which await assignment by the scheduler 603 to a destination microcell. The traffic storage is provided by a switch 609, employing separate (logical) queues 605 that correspond to the individual microcells. It is recognized by one of ordinary skill in the art that any type of switches can be used; e.g., cell-based switches and frame-based switches. That is, the queues 605 store the traffic that are received from the transmitting STs (not shown). The scheduler 603 serves the queues 607 of the switch 609. Within the satellite 601, a Bandwidth Control Processor (BCP) 611, as previously discussed, regulates the granting and denial of bandwidth requests from the transmitting STs, thereby limiting the amount of traffic that the switch 609 processes.

[0052] During busy periods for the system 600, the traffic demand to some microcells may greatly exceed the capacity of the microcell to receive traffic. Without traffic control to mitigate this situation, the queues 607 for such microcells would quickly fill and overflow, causing the switch 609 to become congested for those microcells and to drop packets. This congested state of the switch 609 negatively impacts system performance.

[0053] Thus, it is highly desirable that such switch congestion be avoided, and that the transmitting STs be given the opportunity to shed excess traffic gracefully. Another concern is that the amount of bandwidth that is available for potentially congested microcells is likely to vary dynamically due to competition and potential interference with other microcells.

[0054] Consequently, the satellite communication system 600 employs a traffic control processing logic 613. The interaction between the traffic control processing logic 613, the scheduler 603, and the BCP 611 constitutes an Available Bandwidth Control mechanism (ABCM) that can dynamically adapt to traffic conditions to avoid congestion within the satellite communication system 600. As shown in Fig. 6, the traffic control processing logic 613 is located external to the satellite 601; for example, the traffic control processing logic 613 may reside in a NOC 507 (Figure 5). In a system

in which processing resources are a concern, implementing the traffic control processing logic 613 at a site remote from the satellite 601 would be preferable, so that the processing and power resources of satellite 601 can be utilized for other functions. However, depending on the specific application of the satellite communication system 600, the traffic control processing logic 613 may alternatively be housed within the satellite 601. In other words, the traffic control processing logic functions may be placed within the satellite 601 or within a ground-based center (e.g., NOC).

[0055] For explanatory purposes, Figure 6 illustrates an embodiment in which the traffic control processing logic 613 is placed in the NOC 507 (Figure 5). The ABCM, of which the traffic control processing logic 613 is a key part effectively provides a real-time traffic control loop. The traffic control processing logic 613 includes a microcell utilization module 615 for determining utilization, an error calculation module 617 for comparing the computed utilization with a target utilization, a gain and filtering module 619 for computing a correction value, an adder 621 for outputting a control value, and a control storage 623 for storing the control values. It is known to the art that congestion avoidance operation may be obtained by controlling traffic admittance to the network such that the determined utilization is less than 100%, and herein specifically that it should approximate the designated target utilization. The operation of the traffic control processing logic 613 is more fully described with respect to Figure 7.

[0056] Figure 7 shows a flow chart of the operation of the available bandwidth control mechanism (ABCM) of Figure 6. For a given measurement period, the scheduler 603 monitors the traffic that is assigned for transmission to a microcell, gathering the metrics associated with the observed traffic that is processed from the queues 607. In an exemplary embodiment, the metrics are the counter values of the Yes counter and the No counter, as described above with respect to Figure 5. The counters are incremented from the start of a measurement period until the period expires; the measurement period can be set to any prescribed time interval, depending on the desired response of the traffic control processing logic 613. At the completion of a measurement period, the scheduler 603, as in step 701, provides the collected bandwidth metrics on each microcell of interest to the traffic control processing logic 613. These metrics are formulated in a message and sent to the traffic control processing logic 613 by the satellite 601, where they are received by a utilization module 615. In step 703, the utilization module 615 determines the utilization of each microcell. In particular, based on the received bandwidth metrics, the utilization module 615 calculates utilization for each microcell as follows:

$$\text{Utilization} = \text{Carried} / (\text{Carried} + \text{Additional}), \quad \text{Eq. (5)}$$

where Carried represents the actual traffic transmitted to a microcell and Additional represents available but unused bandwidth. The units of the Carried and Additional parameters may be any convenient, consistent units; e.g., megabits-per-second, kilo-packets-per-second (kpps), or bursts-per-millisecond.

[0057] It is recognized by one of ordinary skill in the art that Eq. 5 and the ABCM equations based upon Eq. (5) have general utility. When embodied into a system, such as system 500, which utilizes a scheduler 513, providing Yes and No counts, the specific implementation of Eq. (5) reduces to that of Eq. (4).

[0058] For each microcell, there is a target utilization, which is a parameter that may be fixed or may be adaptive to general traffic conditions. To avoid congestion and possible packet loss, the utilization for any microcell should be less than 100%. Accordingly, the target utilization may initially be set to a default value of, for example, 95%, and may subsequently be tailored to reflect the traffic pattern corresponding to the particular microcell. In step 705, the error calculation module 617 computes the difference (or error) between the computed utilization and a corresponding target utilization for each microcell according to the following equation:

$$\text{Error} = \text{Target} - \text{Utilization}, \quad \text{Eq. (6)}$$

where Error and Target are dimensionless fractions. Alternatively, the Error may be expressed in traffic units, as follows:

$$\text{Error} = (\text{Target}/100) * (\text{Carried} + \text{Additional}) - \text{Carried} \quad \text{Eq. (7)}$$

[0059] For computational efficiency, the above Error value is formulated to exhibit the same polarity as that of the Correction value (i.e., all positive values or all negative values). Next, in step 707, the error calculation module 617 determines whether or not the difference between the computed utilization and the corresponding target utilization is within a predetermined range (e.g., 3% variation). If the utilization of a particular microcell matches the Target value or falls within the predetermined range, then zero error is determined. If the difference falls outside the predetermined range, then the Correction value, as in step 709, is determined. For each microcell, any significant error is passed on to the gain and filtering module 619, which determines the correction value.

[0060] According to an exemplary embodiment, the difference represents the correction value (or error value); the direct use of the Error value as the Correction value produces a unity loop gain.

[0061] Although the Error value may be directly used as the Correction value, it is often desirable to employ additional processing. The bandwidth metric values from the scheduler 603 may exhibit noisy fluctuations, especially if the measurement periods are of short durations. In this case, it may be desirable for module 619 to employ mechanisms for filtering and/or smoothing of the measurement samples.

[0062] Additionally, it may be desired to deliberately slow the response of the traffic control loop of the ABCM, in which case the gain and filtering module 619 of the traffic control processing logic 613 may employ an algorithm gain that is less than unity. Alternatively, it may be desirable to achieve a fast-attack to correct utilization excesses and/or achieve a slow-increase of the carried traffic when the utilization is nearing the target value. In such a case, split gain may be used. In addition, the gain and filtering module 619 provides reconciliation of any differences in the operating intervals or units employed between the BCP 611 and the scheduler 603.

[0063] Thereafter, the Control Value is computed from the Correction value. The output of the gain and filtering module 619 (which is the Correction value) is input to an adder 621.

The Control Value that is output from the adder 621 is based on the following:

$$\text{New Control Value} = \text{Reference Control Value} + \text{Correction Value.} \quad \text{Eq.(8)}$$

The Reference Control Value is the Control Value that was previously transmitted to the BCP 611 during the measurement period that produced the metrics that were used in deriving the current Correction value. A key value for facilitating this correlation is that of a Round-Trip Control Time (RTCT). This RTCT is the time between successive, causal adjustments in the Control Values. The RTCT includes all causal delays in one round trip of the control loop; e.g., propagation delays, message transmission/reception delays, message processing delays, the measurement period, etc.

[0064] For Reference Control Value lookup purposes, in one embodiment of the present invention, the RTCT value is converted into an index value that represents how many Control Value messages ago that the Reference Control Value was sent out. In the case in which the Control Value messages are sent out at regular intervals, and for which $\text{RTCT}_{\text{seconds}}$ is known, the index is computed as follows:

$$\text{RTCT}_{\text{Index}} = \text{Int.} [\text{RTCT}_{\text{seconds}} / \text{Interval}], \quad \text{Eq.(9)}$$

where Int. [] represents the value of the division, wherein any fractional part is rounded up to the next integer. For example, if the RTCT is 1.7 seconds, and the message interval is 0.5 seconds, then $\text{RTCT}_{\text{Index}}$ equals 4.

[0065] However, if the RTCT is not constant or is otherwise uncertain, then an adaptive approach is used to obtain either the Reference Control Values or $\text{RTCT}_{\text{Index}}$. The Reference Control Values may be obtained from the BCP 611 by reading those control values that were active during the measurement period. The BCP 611 sends the appropriate prior control values to the traffic control processing logic 613 along with the current bandwidth metrics. Under this approach, the control storage 623 is not required. This approach is particularly effective when the traffic control processing logic 613 is collocated with the scheduler 603. If the traffic control processing logic 613 is not collocated with the scheduler 603 (e.g., ground based), then storing the values locally within the traffic control processing logic 613 may be preferable, as this may represent a significant amount of control information.

[0066] In step 711, the Correction value is applied to a selected, previously stored Control Value for the microcell to produce a new Control Value for the particular microcell. On the other hand, if the difference falls within the predetermined range that corresponds to zero error, then the new Control Value is set to the previous Control Value (step 713). The New Control Value is stored for later use in control storage 623, per step 715.

[0067] According to an embodiment of the present invention, the control storage index can be dynamically determined without the requirement that the RTCT be known or constant. As the gain and filtering module 619 produces the Correction values, the module 619 also produces Control Sequence Numbers (CSNs) to correspond to the new Control Values. The CSNs follow a modulo-n function (i.e., the numbering is cyclic). The CSN is stored in the control storage 623 and is sent to the BCP 611, which also stores the CSN locally. When the bandwidth metrics are determined by scheduler 603, scheduler 603 also reads the value of the CSN that is associated with the measurement period. Scheduler 603 includes the CSN with the metrics in the message that is transmitted to the traffic control processing logic 613. Thus, when the Correction values are determined from these metrics, as in step 709, the respective CSNs are employed as indices into a table within the control storage 623, yielding the proper Reference Control Values, to be employed by the adder 621.

[0068] In step 717, the Control Value is sent within a message to the BCP 611 in the satellite 601. Accordingly, the

BCP 611 regulates the uplink traffic from the transmitting STs, per step 719, such that, during any one bandwidth allocation period, the traffic authorized by the BCP 611 for any particular microcell does not exceed the limit established by the traffic control processing logic 613 for that microcell. It is noted that the actual traffic admitted to the switch 609 and stored in the queue 607 for a particular microcell may be less than the limit authorized by the traffic control processing logic 613 if the aggregate transmitting ST demand to the microcell is less than this limit (as specified by the Control Value).

[0069] The scheduler 603 responds adaptively, assigning bandwidth as may be available to those microcells whose queues contain traffic. For some microcells, the scheduler 603 may have additional bandwidth available beyond that of the actual carried traffic for the microcell. The monitoring function of the scheduler 603 collects and reports, at least, the carried and available bandwidth for each microcell of interest, initiating a repeat of the steps 701-719 for each such microcell.

[0070] As part of a traffic control loop, the traffic control processing logic 613 is stable at unity gain for which, in response to a transient, the logic 613 exhibits neither overshoot nor undershoot. It is observed that a causal correction cannot be made faster than one RTCT. The traffic control processing logic 613 advantageously provides 100% of the indicated correction in exactly one RTCT with stable loop operation. This characteristic of the present invention provides a major improvement over conventional systems. Given its adaptive capability, the traffic control processing logic 613 is well-suited for bursty traffic whose bandwidth needs are not known in advance. That is, no prior assumptions as to traffic profiles or patterns are required for the operation of the traffic control processing logic 613. Consequently, the traffic control processing logic 613 is also adaptive to the presence of allocated traffic. This transparency to the traffic pattern advantageously enhances operation of communication system 600, by significantly reducing, for example, the operational load on a system operator.

[0071] To better appreciate the present invention, it is instructive to analyze the performance of the ABCM. The transient performance of the closed-loop traffic control algorithm of Figure 6 significantly impacts the system performance. The transient performance is characterized by a combination of the speed of response, post-transient stability including lack of any ringing, and accuracy in maintaining the desired bandwidth limit when subjected to varying degrees of available bandwidth and of the actual traffic due to any traffic allocations. If the loop responds too slowly, then any sudden decrease in bandwidth available to a queue 605 may result in an unacceptable amount of packet loss stemming from queue overflow. Conversely, if additional bandwidth suddenly becomes available, and the loop is slow to react, then this temporarily available bandwidth may be lost to the system 600. This dynamic behavior is largely dependent on the gain and reference mechanisms employed, the control loop's RTCT, and the values for the several gain parameters to be employed.

[0072] The overall loop gain is the product of two gains: the algorithm gain and a message gain. For the system of Figure 6, the algorithm gain is that of module 619. One simple, but effective implementation of the gain function for module 619 is to relate the Error of Eq. (7) with the Correction Value of Eq. (9) by a simple linear factor:

$$\text{Correction Value} = \text{Gain} * \text{Error} \quad \text{Eq.(9)}$$

This Gain parameter basically determines the speed with which the bandwidth control limits are updated in response to a transient. Conceptually, a simple, but effective form of bandwidth adjustment is obtained by combining Eq. (9) with Eq. (7) to obtain:

$$\text{New Control Value} = \text{Reference Control Value} + \text{Gain} * \text{Error} \quad \text{Eq.(10)}$$

Ideally, one would like the Gain of Eq. (10) to be unity (i.e., 1.00). With unity gain, when an error is observed, the Control Values are corrected promptly and exactly (within measurement errors). However, unless a proper selection is made for the Reference Control Value, the post-transient response can be badly impaired, and even become unstable, with unity Gain due to the message gain effect. Conceptually, the message gain is an integer that equals the number of outstanding messages in the traffic control loop rounded up to the nearest whole message. This gain effect is dependent on both the RTCT and the message update interval for the bandwidth control loop of a system such as 600. For ground-based control of satellite bandwidth, RTCT is typically about 1.5 to 2 seconds, and message intervals typically range from 250 to 750 msec., yielding message gains of typically 3 to 6. However, this message gain effect can be nullified by appropriate implementation of the teachings of this invention, in particular, the employment of an appropriate selection for the Reference Control Value as with ABCM.

[0073] The ABCM mechanism avoids the message gain effect and provides stability of the traffic control loop by correlating the appropriate Reference Control Values (RCVs) with the respective measurement periods. ABCM makes direct, immediate bandwidth changes and thereby provides for an ideal transient response. In other words, it is desirable

to make an immediate correction to any indicated changes, in which 100% of the indicated correction (error) is made in one step, and at one RTCT. The ABCM mechanism is further described, as follows:

$$\text{New Control Value} = \text{RCV} + \text{Correction Value} = \text{RCV}_j + \text{Function}(\text{Measurement}_j) \quad \text{Eq. (11)}$$

where, j designates some particular RCV and the bandwidth measurements resulting from this RCV. As evident from Eq. (11), the RCV that is to be used is not that of the previous value of New Control Value, but rather that of the New Control Value that was active when the measurement was made. This correlation may be facilitated by means such as Control Storage 623. Under this approach, the Gain parameter can be set to 100% with complete closed-loop stability. As a result, 100% of the indicated correction can be made in a single correction step, with zero overshoot and zero undershoot - thus obtaining an ideal transient response. Furthermore, Eq. (11) is independent of the RTCT.

[0074] Figure 8 shows a diagram of the scheduler of Figure 5 utilizing four counters for collecting the bandwidth metrics, in accordance with an embodiment of the present invention. Scheduler 603 includes a counter logic 801 that controls the incrementing and decrementing of a Skip counter 803, a Block counter 805, a Yes counter 807, and a No counter 809. These scheduler counters 803, 805, 807, and 809 provide a mechanism to collect performance metrics.

[0075] Several reasons exist to monitor the performance of the scheduler 603. A first metric is to measure periodically the amount of traffic actually assigned to each microcell (i.e., throughput). This throughput measurement, as will be more fully described later, can be determined at the satellite 505 or by separate ground-based STs 501 and 503 in each microcell.

[0076] This throughput metric by itself provides insufficient information with respect to scheduler efficiency, which indicates the amount of the traffic demand that is capable of being assigned - considering all of the real-time restrictions - is actually being assigned. Additionally, each mechanism of the scheduler 603 that enforces a restriction may have some degree of adjustability in meeting the requirement. Making such adjustments to the scheduler 603 mechanisms in an optimum manner requires additional information (i.e., metrics) as to the operation of the scheduler 603 - on a microcell-by-microcell basis. In particular, information relating to ingress traffic to the queues 607 is helpful in preventing queue overflow. If congestion arises at the queues 607, overflow will result in lost traffic for the microcells corresponding to the overflowed queue. Accordingly, control of the traffic to some or all such microcells is needed by determining the amount of bandwidth that is available for assignment to the microcells, not just the actual traffic assigned. For example, the amount of traffic that is controlled at the source results in a particular microcell receiving 50 Mbps, as regulated by the scheduler 603.

[0077] The throughput to the particular microcell may be dictated by the following scenarios. Under the first scenario, scheduler 603 has 52 Mbps available, with 50 Mbps being currently used. This situation indicates a normal operating point - approximately 96% utilization of available bandwidth; thus, no corrective action needs to be taken. In a second possible scenario, the scheduler 603 has 70 Mbps available, but only 50 Mbps is currently being demanded and used. In a third scenario, scheduler 603 has 70 Mbps available, in which only 50 Mbps is being currently used. However, the traffic demand is for 65 Mbps. In such a case, the scheduler 603 needs to make adjustments to increase utilization via the ABCM. To make these necessary adjustments, additional metrics relating to the operation of the scheduler 603 is required.

[0078] Figure 9 shows a flowchart of the scheduler monitoring operation, according to an embodiment of the present invention. At the beginning of each measurement period, the counter values associated with a particular queue (microcell) are either noted or are initialized to a known value (e.g., zero), per step 901. The scheduler 603 monitors its own operation by utilizing the four counters 803, 805, 807, and 809 for each microcell (that requires monitoring). The counter logic 801 monitors the state of each of the subject queues 607, and increments the appropriate counters 803, 805, 807, and 809 (per step 903). Next, the control logic 801, as in step 905, determines whether the measurement period has expired. If the measurement period has not expired, then step 903 is repeated until the measurement period does expire. Upon expiration of the measurement period, the traffic control processing logic 613 retrieves the counter values, per step 907, and computes the scheduler performance metrics (step 909). The operation of control logic 801 is discussed more fully in Figure 10.

[0079] Figure 10 is a flowchart of the metric collection process using the counters of Figure 8, according to an embodiment of the present invention. In step 1001, control logic 801 determines whether the queue 607 has been skipped by the scheduler 603. The queue 607 may be skipped for a number of reasons, such as implementation of a policy. The control logic 801 increments the Skip counter 803 in this instance (step 1003) and exits. Else, as in step 1005 the control logic 801 determines whether the queue 607 is blocked by a previous assignment to one or more other microcell(s); if so, control logic 801, as in step 1007, increments the Block counter 805 and exits. Else, as in step 1009, control logic 801 checks whether the queue 607 is ready for assignment. If queue 607 is not ready, then the No counter 807 is incremented, per step 1011, and exits. That is, the queue 607 has nothing to send, and thus, the scheduler 603 has made no assignment. If the queue 607 is ready for assignment (as determined by control logic 801), then the Yes

counter 809 is incremented (step 1013).

[0080] Thereafter, at some measurement period. Time, the counters 803, 805, 807, and 809, are read for each microcell (step 907), and another measurement interval is started. These measurement results (metrics) are sent to one (or more) facilities to evaluate the scheduler 603 performance.

[0081] A variety of derived scheduler 603 performance metrics may be obtained from the four metrics corresponding to the four counters 803, 805, 807, and 809.

$$\text{Throughput} = \text{Yes} / \text{Time} \quad \text{Eq. (12)}$$

$$\text{Extra bandwidth available} = \text{No} / \text{Time} \quad \text{Eq. (13)}$$

$$\text{Total bandwidth available} = (\text{Yes} + \text{No}) / \text{Time} \quad \text{Eq. (14)}$$

$$\text{Percent utilization of available bandwidth} = 100\% * \text{Yes} / (\text{Yes} + \text{No}) \quad \text{Eq. (15)}$$

$$\text{Scheduling rate} = (\text{Block} + \text{Yes} + \text{No}) / \text{Time} \quad \text{Eq. (16)}$$

$$\text{Percent blocking} = 100\% * \text{Block} / (\text{Block} + \text{Yes} + \text{No}) \quad \text{Eq. (17)}$$

$$\text{Percent skipping} = 100\% * \text{Skip} / \text{Total} \quad \text{Eq. (18)}$$

where

$$\text{Total} = \text{Skip} + \text{Block} + \text{Yes} + \text{No}.$$

[0082] Depending on the adjustment and control capabilities desired, not all of the four counters 803, 805, 807, and 809 (per microcell) need to be implemented in a specific design. As indicated by Eq. (12), implementation of the Yes counters 809 provides the basic throughput metric (which is also available by ground-based mechanisms). All of the remaining secondary metrics are very difficult or impractical to achieve with ground-based mechanisms. The addition of the No counter 807 (per microcell) permits determination of the effectiveness of the traffic control measures taken, and may also serve as a real-time input metric for the ABCM, using, for example, Eqs. (13) - (15).

[0083] Eq. (15) provides the basis for design of a mechanism for determining the actual utilization of a queue in terms of its egress traffic. According to one embodiment of the present invention, enhanced operation may be obtained by modifying Eq. (15), as described below. Under certain circumstances, in which there is either a large, abrupt increase in traffic that is offered to the queue and/or a large abrupt, decrease in the bandwidth available for the egress traffic (i. e., a sudden drop in the queue exit traffic rate), Eq. (15) can be altered to account for such conditions. For the purposes of illustration of the scenario involving an abrupt decrease in the available bandwidth of system 600, the following assumptions are made: the Target utilization value is 95%, the available bandwidth is 200 kpps, and the queue is in steady state whereby the actual utilization, as given by Eq. (15), equals the Target value set for it (i.e., the Error is zero). Given a queue input of 190 kpps, a sub-system failure may cause the available bandwidth to abruptly drop to 140 kpps. As a result, the actual queue utilization rises to 100%, with an indicated percentage Error of - 5% (95% - 100%). Using Eq. (7), the Error becomes - 7.6 kpps; i.e., $(95/100) * (140 + 0) - (140)$, leading to a New Control Value of 182.4 kpps $(190 - 7.6)$. This adjustment may be inadequate under this circumstance as the desired control value for 95% utilization is 133 kpps $(95/100 * 140)$. This potential deficiency stems from the fact that the measured utilization is capped out at 100%. To address this situation, the concept of an "offered" utilization is defined as follows:

$$\text{Offered Utilization} = 100\% * \text{Offered_to_queue} / (\text{Yes} + \text{No}) \quad \text{Eq. (19)}$$

Introducing this concept into Eq. (7) yields the following Error calculation:

$$\text{Error} = (\text{Target}/100) * (\text{Yes} + \text{No}) - \text{Offered_to_queue} \quad \text{Eq.(20)}$$

Continuing with the above example, using Eq. (20) rather than Eq. (7), the Error is - 57 kpps (i.e., $(95/100) * (140 + 0) - (190)$). This new Error calculation provides a New Control Value of 133 $(190 - 57)$, which is the correct value.

[0084] For cases in which the Offered_to_queue is unknown, an approximation to Eq. (20) may be obtained by replacing Offered_to_queue with $(\text{Yes} + \text{Loss})$, where Loss is the input loss rate to the queue due to queue overflow, assuming that Loss is in the same units as Yes and No:

$$\text{Error} = (\text{Target}/100) * (\text{Yes} + \text{No}) - (\text{Yes} + \text{Loss}) \quad \text{Eq.(21)}$$

During operation at high load but without congestion, the Loss will be zero and No value will be small (but non-zero). By contrast, in the case in which an abrupt and large decrease in available bandwidth is experienced, the No variable is zero and Loss is a significant value.

[0085] Module 613 may employ mechanisms for Eqs. (19), (20) and/or (21) for the situations in which higher performance of system 600 is required. In such a case, module 609 may supply the Offered_to_queue or Loss metrics to 613, in addition to the Yes and No metrics that are furnished by module 603.

[0086] A measurement of system fairness may be obtained by comparing the results of Eq. (14) for one microcell with that of the other microcells. However, any resulting unfairness may be the result of the (necessary) assignment restrictions that the scheduler 603 may need to exert. For this situation, a measurement of scheduler fairness may be obtained by incorporating the Block counter 805 and comparing the results of Eq. (16) for one microcell with that of the others. If the results of Eq. (17) appear to be too high for certain microcells, then the particular scheduler 603 algorithms may require adjustment, particularly if other microcells are receiving too little traffic.

[0087] Inclusion of a Skip counter 803 or equivalent metric is helpful in ascertaining the results and effectiveness of any deliberate weighting functions that may be employed in the scheduling algorithms.

[0088] Extensions to the above metrics may be utilized, depending on the applications. For example, additional counters may be employed for diagnostic purposes; a number of counters could be designated to explicitly indicate the various causes of blocking. It should be noted that the measurement period could be designated differently for different counters. If real-time control is not required with respect to blocking and skipping, then the interval at which these are metrics are measured can be significantly increased compared to those counters that are used for real time measurements (e.g., Yes and No counters 809 and 807). In this case, suitable time aggregation of the blocking and skipping counter values would be employed for use in Eqs. (16) - (18).

[0089] Figure 11 illustrates a computer system 1101 upon which an embodiment according to the present invention may be implemented to perform congestion avoidance. Computer system 1101 includes a bus 1103 or other communication mechanism for communicating information, and a processor 1105 coupled with bus 1103 for processing the information. Computer system 1101 also includes a main memory 1107, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 1103 for storing information and instructions to be executed by processor 1105. In addition, main memory 1107 may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 1105. Computer system 1101 further includes a read only memory (ROM) 1109 or other static storage device coupled to bus 1103 for storing static information and instructions for processor 1105. A storage device 1111, such as a magnetic disk or optical disk, is provided and coupled to bus 1103 for storing information and instructions.

[0090] Computer system 1101 may be coupled via bus 1103 to a display 1113, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 1115, including alphanumeric and other keys, is coupled to bus 1103 for communicating information and command selections to processor 1105. Another type of user input device is cursor control 1117, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 1105 and for controlling cursor movement on display 1113.

[0091] According to one embodiment, the traffic control processing logic 613 (Figure 6) is performed by computer system 1101 in response to processor 1105 executing one or more sequences of one or more instructions contained in main memory 1107. Such instructions may be read into main memory 1107 from another computer-readable medium, such as storage device 1111. Execution of the sequences of instructions contained in main memory 1107 causes processor 1105 to perform the process steps described herein. One or more processors in a multi-processing arrangement may also be employed to execute the sequences of instructions contained in main memory 1107. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions. Thus, embodiments are not limited to any specific combination of hardware circuitry and software.

[0092] Further, the congestion avoidance mechanism of the present invention may reside on a computer-readable medium. The term "computer-readable medium" as used herein refers to any medium that participates in providing

instructions to processor 1105 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 1111. Volatile media includes dynamic memory, such as main memory 1107. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 1103. Transmission media can also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communication.

[0093] Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0094] Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 1105 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions relating to the notification services to control call processing remotely into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 1101 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector coupled to bus 1103 can receive the data carried in the infrared signal and place the data on bus 1103. Bus 1103 carries the data to main memory 1107, from which processor 1105 retrieves and executes the instructions. The instructions received by main memory 1107 may optionally be stored on storage device 1111 either before or after execution by processor 1105.

[0095] Computer system 1101 also includes a communication interface 1119 coupled to bus 1103. Communication interface 1119 provides a two-way data communication coupling to a network link 1121 that is connected to a local network 1123. For example, communication interface 1119 may be a network interface card to attach to any packet switched local area network (LAN). As another example, communication interface 1119 may be an asymmetrical digital subscriber line (ADSL) card, an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. Wireless links may also be implemented. In any such implementation, communication interface 1119 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0096] Network link 1121 typically provides data communication through one or more networks to other data devices. For example, network link 1121 may provide a connection through local network 1123 to a host computer 1125 or to data equipment operated by a service provider, which provides data communication services through a communication network 1127 (e.g., the Internet). LAN 1123 and network 1127 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 1121 and through communication interface 1119, which carry the digital data to and from computer system 1101, are exemplary forms of carrier waves transporting the information. Computer system 1101 can transmit notifications and receive data, including program code, through the network(s), network link 1121 and communication interface 1119.

[0097] The techniques described herein provide several advantages over prior approaches to avoiding traffic congestion. A scheduler 613 generates bandwidth metrics for a destination region (e.g., microcell). A traffic control processing logic 613 receives the bandwidth metrics. The traffic control processing logic 613 includes a utilization module 615 that determines utilization associated with the destination site based upon the received bandwidth metrics, and an error calculation module 617 that computes the difference between the determined utilization and a target utilization. A gain and filtering module 619, which is also a part of the traffic control processing logic 613, computes a Correction value based upon the difference between the determined utilization and the target utilization, in which the Correction value being associated with the destination site. Further, the traffic control processing logic 613 includes an adder 621 that outputs a Control Value based upon a Reference Control Value and the Correction value. A bandwidth control processor 611 assigns bandwidth allocation based upon the control value. This approach advantageously provides enhanced system efficiency. Another advantage is that this arrangement optimizes system throughput.

[0098] Obviously, numerous modifications and variations of the present invention are possible in light of the above teachings. It is therefore to be understood that within the scope of the appended claims, the invention may be practiced otherwise than as specifically described herein.

Claims

1. A method of measuring scheduler efficiency, the method comprising:

selectively incrementing a plurality of scheduler counters (803, 805, 807, and 809) corresponding to a plurality of scheduling events; and
computing performance metrics based upon values of the plurality of scheduler counters (803, 805, 807, and

809).

2. The method according to Claim 1, wherein the performance metrics in the computing step comprises at least one of throughput, extra bandwidth available, total bandwidth available, percent utilization of available bandwidth, scheduling rate, percent blocking, and percent skipping.

3. The method according to Claim 1, further comprising:
sending the counter values to a remote processor (508).

4. The method according to Claim 1, wherein the plurality of scheduler counters (803, 805, 807, and 809) in the step of selectively incrementing comprise a skip assignment counter associated with a scheduling event indicating that the scheduler (603) has skipped assignment for transmission, a blocked assignment counter associated with a scheduling event indicating that the scheduler (603) has blocked assignment for transmission, a No counter (807) associated with a scheduling event indicating that a queue (607) is not ready for assignment, and a Yes counter (809) associated with a scheduling event indicating that the queue (607) is ready for assignment.

5. The method according to Claim 4, wherein the computing step comprises:

computing throughput based upon a value of the Yes counter (809);
computing extra bandwidth available based upon a value of the No counter (807);
computing total bandwidth available based upon the value of the Yes counter (809) and the value of the No counter (807);
computing percent utilization of available bandwidth based upon the value of the Yes counter (809) and the value of the No counter (807);
computing scheduling rate based upon a value of the block counter (805), the value of the Yes counter (809), and the value of the No counter (807);
computing percent blocking based upon the value of the block counter (805), the value of the Yes counter (809), and the value of the No counter (807); and
computing percent skipping based upon a value of the skip counter (803).

6. The method according to Claim 4, wherein the computing step comprises:
computing percent of offered utilization of available bandwidth based upon the value of the Yes counter (809), the value of the No counter (807), and an Offered-to-Queue value that reflects amount of traffic from a plurality of queues (607).

7. The method according to Claim 4, wherein the scheduler (603) regulates transmission to a microcell, the skip counter (803), the block counter (805), the No counter (807), and the Yes counter (809) being associated with the microcell.

8. The method according to Claim 1, wherein the plurality of scheduler counters (803, 805, 807, and 809) in the step of selectively incrementing include diagnostic counters to indicate causes of blocking.

9. A scheduling device having metrics measurement capability, comprising:

a plurality of scheduler counters (803, 805, 807, and 809) corresponding to a plurality of scheduling events; and
a counter logic (801) configured to selectively increment the plurality of scheduler counters (803, 805, 807, and 809), the counter logic (801) being coupled to a plurality of queues (607) that store cells,

wherein performance metrics are computed based upon values of the plurality of scheduler counters (803, 805, 807, and 809).

10. The device according to Claim 9, wherein the performance metrics comprises at least one of throughput, extra bandwidth available, total bandwidth available, percent utilization of available bandwidth, scheduling rate, percent blocking, and percent skipping.

11. The device according to Claim 9, wherein the counter values are transmitted to a remote processor (508).

12. The device according to Claim 9, wherein the plurality of scheduler counters (803, 805, 807, and 809) comprise

a skip assignment counter associated with a scheduling event indicating that the scheduler (603) has skipped assignment for transmission, a block counter (805) associated with a scheduling event indicating that the scheduler (603) has blocked assignment for transmission, a No counter (807) associated with a scheduling event indicating that a queue (607) is empty, and a Yes counter (809) associated with a scheduling event indicating that the queue (607) is not empty.

13. The device according to Claim 12, wherein the performance metrics comprises throughput that is computed based upon a value of the Yes counter (809); extra bandwidth available that is computed based upon a value of the No counter (807); total bandwidth available that is computed based upon the value of the Yes counter (809) and the value of the No counter (807); percent utilization of available bandwidth that is computed bandwidth based upon the value of the Yes counter (809) and the value of the No counter (807); scheduling rate that is computed based upon a value of the block counter (805), the value of the Yes counter (809), and the value of the No counter (807); percent blocking that is computed blocking based upon the value of the block counter (805), the value of the Yes counter (809), and the value of the No counter (807); and percent skipping that is computed based upon a value of the skip counter (803).
14. The device according to Claim 12, wherein the performance metrics comprises percent of offered utilization of available bandwidth based upon the value of the Yes counter (809), the value of the No counter (807), and an Offered-to-Queue value that reflects amount of traffic from the plurality of queues (607).
15. The device according to Claim 12, wherein the scheduling device regulates transmission to a microcell, the skip counter (803), the block counter (805), the No counter (807), and the Yes counter (809) being associated with the microcell.
16. The device according to Claim 9, wherein the plurality of scheduler counters (803, 805, 807, and 809) include diagnostic counters to indicate causes of blocking.
17. A switching system comprising:
 - a plurality of input ports (101) configured to receive cells;
 - a plurality of output ports (113) communicating with a plurality of destination sites; and
 - a plurality of queues (607) coupled to the input ports (101) and the output ports (113) for storing the received cells; and
 - a scheduler (603) coupled to the plurality of queues (607) and configured to assign the stored cells for transmission out the output ports (113), the scheduler (603) comprising,
 - a plurality of scheduler counters (803, 805, 807, and 809) corresponding to a plurality of scheduling events, and
 - a counter logic (801) configured to selectively increment the plurality of scheduler counters (803, 805, 807, and 809); and
 - a processor (508) configured to compute performance metrics based upon values of the plurality of scheduler counters (803, 805, 807, and 809).
18. The system according to Claim 17, wherein the performance metrics comprises at least one of throughput, extra bandwidth available, total bandwidth available, percent utilization of available bandwidth, scheduling rate, percent blocking, and percent skipping.
19. The system according to Claim 17, wherein the processor (508) is located remotely from the scheduler (603).
20. The system according to Claim 17, wherein the plurality of scheduler counters (803, 805, 807, and 809) comprise a skip assignment counter associated with a scheduling event indicating that the scheduler (603) has skipped assignment for transmission, a block counter (805) associated with a scheduling event indicating that the scheduler (603) has blocked assignment for transmission, a No counter (807) associated with a scheduling event indicating that a queue (607) is empty, and a Yes counter (809) associated with a scheduling event indicating that the queue (607) is not empty.
21. The system according to Claim 20, wherein the performance metrics comprises throughput that is computed based upon a value of the Yes counter (809); extra bandwidth available that is computed based upon a value of the No counter (807); total bandwidth available that is computed based upon the value of the Yes counter (809) and the value of the No counter (807); percent utilization of available bandwidth that is computed bandwidth based upon

the value of the Yes counter (809) and the value of the No counter (807); scheduling rate that is computed based upon a value of the block counter (805), the value of the Yes counter (809), and the value of the No counter (807); percent blocking that is computed blocking based upon the value of the block counter (805), the value of the Yes counter (809), and the value of the No counter (807); and percent skipping that is computed based upon a value of the skip counter (803).

22. The system according to Claim 20, wherein the scheduler (603) regulates transmission to a microcell, the skip counter (803), the block counter (805), the No counter (807), and the Yes counter (809) being associated with the microcell.

23. The system according to Claim 20, wherein the performance metrics comprises percent of offered utilization of available bandwidth based upon the value of the Yes counter (809), the value of the No counter (807), and an Offered-to-Queue value that reflects amount of traffic from the plurality of queues (607).

24. The system according to Claim 17, wherein the plurality of scheduler counters (803, 805, 807, and 809) include diagnostic counters to indicate causes of blocking.

25. A computer-readable medium carrying one or more sequences of one or more instructions for measuring scheduler efficiency, the one or more sequences of one or more instructions including instructions which, when executed by one or more processors, cause the one or more processors to perform the steps of:

selectively incrementing a plurality of scheduler counters (803, 805, 807, and 809) corresponding to a plurality of scheduling events; and
computing performance metrics based upon values of the plurality of scheduler counters (803, 805, 807, and 809).

26. The computer readable medium according to Claim 25, wherein the performance metrics in the computing step comprises at least one of throughput, extra bandwidth available, total bandwidth available, percent utilization of available bandwidth, scheduling rate, percent blocking, and percent skipping.

27. The computer readable medium according to Claim 25, further comprising computer-executable instructions for causing the computer system to perform the steps of:
sending the counter values to a remote processor (508).

28. The computer readable medium according to Claim 25, wherein the plurality of scheduler counters (803, 805, 807, and 809) in the step of selectively incrementing comprise a skip assignment counter associated with a scheduling event indicating that the scheduler (603) has skipped assignment for transmission, a block counter (805) associated with a scheduling event indicating that the scheduler (603) has blocked assignment for transmission, a No counter (807) associated with a scheduling event indicating that a queue (607) is empty, and a Yes counter (809) associated with a scheduling event indicating that the queue (607) is not empty.

29. The computer readable medium according to Claim 28, further comprising computer-executable instructions for causing the computer system to perform the steps of:

computing throughput based upon a value of the Yes counter (809);
computing extra bandwidth available based upon a value of the No counter (807);
computing total bandwidth available based upon the value of the Yes counter (809) and the value of the No counter (807);
computing percent utilization of available bandwidth based upon the value of the Yes counter (809) and the value of the No counter (807);
computing scheduling rate based upon a value of the block counter (805), the value of the Yes counter (809), and the value of the No counter (807);
computing percent blocking based upon the value of the block counter (805), the value of the Yes counter (809), and the value of the No counter (807); and
computing percent skipping based upon a value of the skip counter (803).

30. The computer readable medium according to Claim 28, wherein the scheduler (603) regulates transmission to a microcell, the skip counter (803), the block counter (805), the No counter (807), and the Yes counter (809) being

associated with the microcell.

31. The computer readable medium according to Claim 25, wherein the plurality of scheduler counters (803, 805, 807, and 809) in the step of selectively incrementing include diagnostic counters to indicate causes of blocking.

32. The computer readable medium according to Claim 28, wherein the computing step comprises:
computing percent of offered utilization of available bandwidth based upon the value of the Yes counter (809), the value of the No counter (807), and an Offered-to-Queue value that reflects amount of traffic from a plurality of queues (607).




FIG. 1




FIG. 2





FIG. 3A






FIG. 4




FIG. 5




FIG. 6

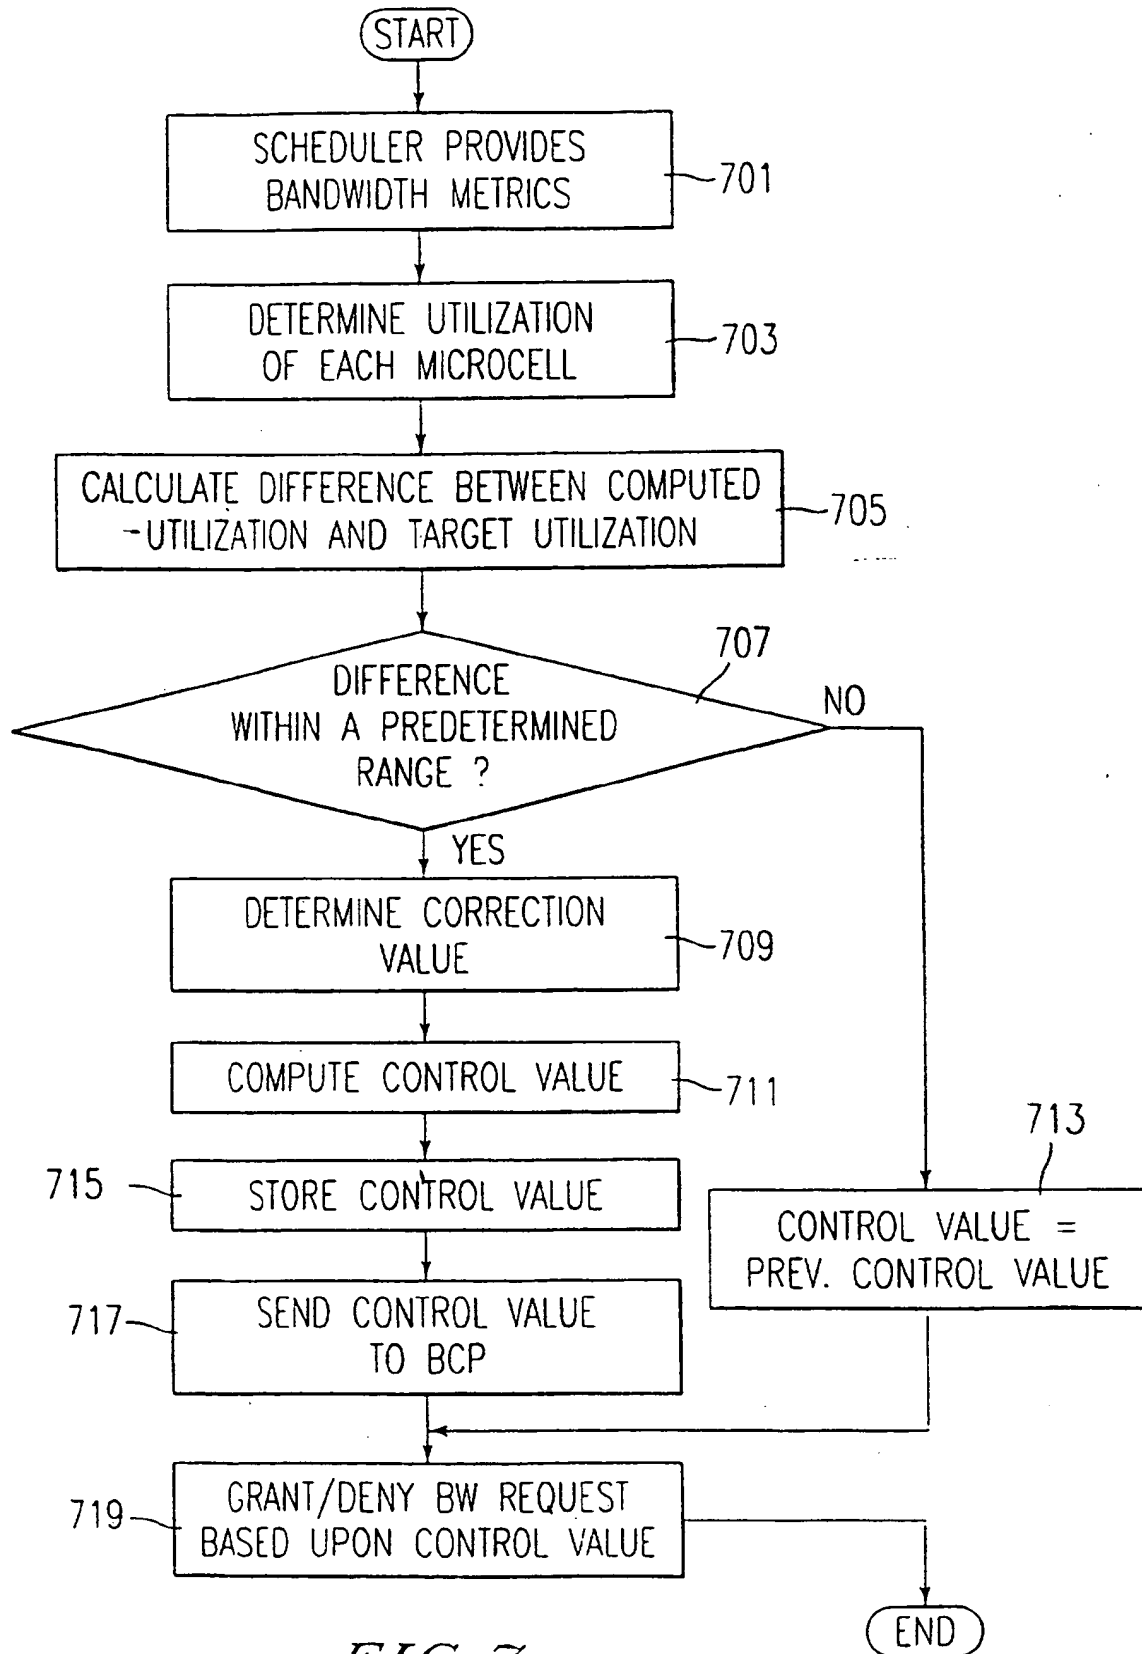


FIG. 7


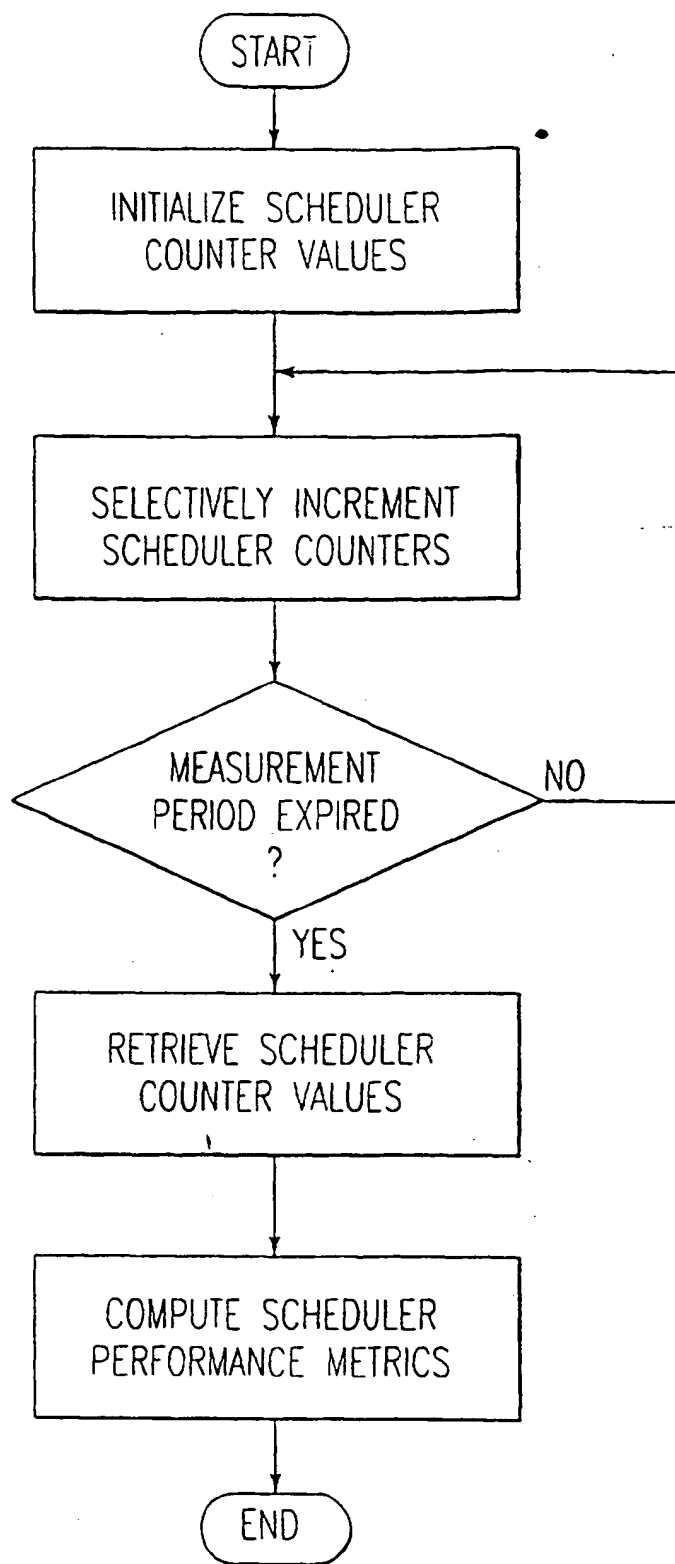


FIG. 8

*FIG. 9*

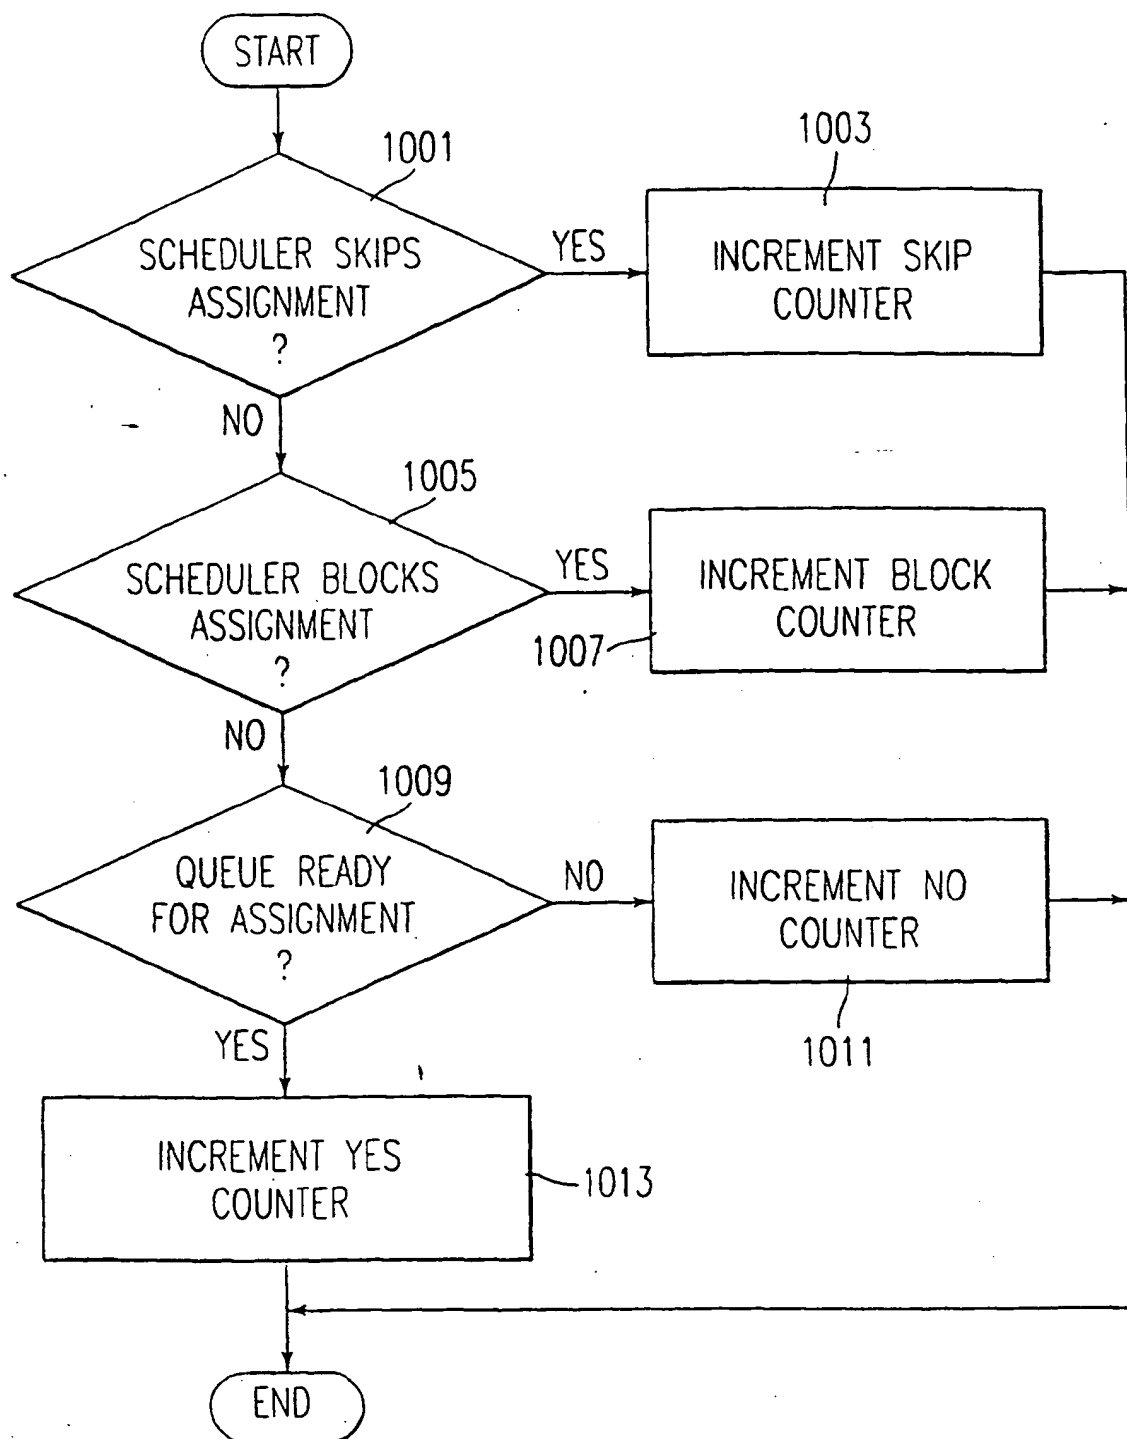


FIG. 10




FIG. 11